

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
9 October 2003 (09.10.2003)

PCT

(10) International Publication Number
WO 03/083656 A2

- (51) International Patent Classification⁷: **G06F 9/48**
- (21) International Application Number: **PCT/IB03/00626**
- (22) International Filing Date: 14 February 2003 (14.02.2003)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
0207296.5 28 March 2002 (28.03.2002) GB
- (71) Applicant (for all designated States except US): **KONINKLIJKE PHILIPS ELECTRONICS N.V. [NL/NL]**; Groenewoudseweg 1, NL-5621 BA Eindhoven (NL).
- (72) Inventor; and
- (75) Inventor/Applicant (for US only): **KING, Colin, I. [GB/GB]**; Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).
- (74) Agent: **WHITE, Andrew, G.**; Internationaal Octrooibureau B.V., Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).
- (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:
— without international search report and to be republished upon receipt of that report
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.



WO 03/083656 A2

(54) Title: METHOD AND APPARATUS FOR CONTEXT SWITCHING IN COMPUTER OPERATING SYSTEMS

(57) Abstract: The present invention provides for a method and apparatus for a computer implemented system including a processor and cache memory arranged to receive operating system instructions for context switching between processes, and including control means for writing back cache data to memory means during processor idle cycle at completion of a process, and prior to initiation of the context switch so as to enhance the operating speed since, at the time of initiating the context switch, the cached data has already been written back to memory.

DESCRIPTION

**METHOD AND APPARATUS FOR CONTEXT SWITCHING
IN COMPUTER OPERATING SYSTEMS**

5

The present invention relates to a method and apparatus for context switching in computer operating systems.

Operating systems are used within computer systems so as to increase
10 the utilisation of the system's processor. In view of the difference in operating speed between the processor and, for example input output devices, the operating system advantageously serves to schedule instructions to the processor in order to limit the time period for which the processor might otherwise remain idle due to its faster operating speed.

15 With developments in multitasking, and multithreading, an increasing number of interruptions are required to be handled by the operating system and correspondingly frequent context switches between different processes are then required.

With, for example, UNIX-like kernels, the switching between running
20 processes should be as fast as possible and employ the minimum possible number of processor cycles. However, before a context switch to a new process can be achieved, the cached data of the previously running process may be written back to memory for consistency between the contents of the cache and memory. The time taken to write the cached data back to memory
25 once it has been identified that a context switch is to be conducted is a disadvantageously limiting feature with regard to the overall time required to execute the context switch.

Attempts have been made to reduce context switching time as illustrated for example in US-A-6 006 320. Here a duplicate arithmetic logic
30 unit, instruction and data cache, register set and load/store unit are provided in an attempt to speed up the context switching process. However, such an approach to this problem can be considered disadvantageous in requiring the

aforementioned duplication of hardware and proves to be unnecessarily complex.

The present invention therefore seeks to provide for a simplified solution
5 to the problem of context switching time delays.

According to one aspect of the present invention there is provided a method of context switching between processes in a computer operating system including writing cached data back to a memory means, comprising the
10 step of the writing cached data back to the memory means during processor idle cycles at completion of a process and prior to initiation of the context switch.

The invention is advantageous in that it employs the idle loop entered by the processor once, for example, a runnable process has completed and
15 while the processor is waiting for the next event, such as the completion of an input/output operation, to occur. During such an idle loop, spare processor cycles are lost and the present invention advantageously makes use of such cycles in order to write-back the cached data to memory.

Thus, at the end of a process, cached data can somewhat automatically
20 be written back to memory so that, should a context switch be required, there is then no need to first write back the cached data to memory before commencing with the context switch. This can increase substantially the switching time required for context switches in an operating system.

The feature of Claim 2 is advantageous in providing for an effective
25 means for indicating that cached data write-back has occurred so as to lead to appropriate subsequent control steps depending on whether a context switch is required or not.

The subject matter of Claims 3-5 advantageously further adapts the subject matter of Claim 2 in a particularly effective manner so as to lead to a
30 suitable form of process switching as required.

According to another aspect of the present invention, there is provided a computer implemented system including a processor and cache memory

arranged to receive operating system instructions for context switching between processes, and including control means for writing back cached data to memory means during processor idle cycles at completion of a process, and prior to initiation of the context switch.

5 In accordance with a further aspect of the present invention, there is provided a computer program product having computer program instructions and arranged for controlling context switching between processes in a computer operating system so as to write cached data back to memory means during processor idle cycles at completion of a process and prior to initiation of
10 the context switch.

As will be appreciated the present invention can therefore be provided as a software, or hardware, solution to the problem discussed above.

The present invention is described further hereinafter, by way of
15 example only, with reference to the accompanying drawings in which:

Fig. 1 is a flow diagram of a context switching operation according to an embodiment of the present invention;

Fig. 2 is a schematic block diagram of a computer system including an operating system and related control means according to an embodiment of
20 the present invention; and

Fig. 3 is a schematic block diagram of a computer system embodying the present invention and illustrating the transfer of cached pages.

As will be appreciated, the invention proposes the use of idling time to
25 write the cached data back to memory of the last running process, hence removing the time penalty of this operation at process switching time. Once the cached data has been written back to memory, a flag for that process is set to indicate that it has been done.

When a switch between runnable processes occurs, the last runnable
30 processes's flag is checked to see if its cached data has been written back during a CPU idle, and if set, the switch does not need to write cache back. However, if the CPU idle has not been entered, and hence the last runnable

processes's flag is not set, then the process switching operating need to write the cache back.

The resulting advantage is that for a system that has free CPU idle time, switching between processes is faster and the system becomes more responsive to the user. With current fast processors, processor idle time is available for most moderately loaded systems. Processor idle time also occurs when processes are using Input/Output steps and the processor is waiting for these transactions to complete.

Turning first to Fig.1, there is provided a flow diagram illustrating an embodiment of the present invention in which a current process 10 is to be rescheduled 12. At step 14 a search for the next runnable process is made and, if one is located, a check at 16 is made to ascertain if the previous running process's flag has been set. If no flag has been set then a decision is made at 18 to copy back the cache of the previous running process and the flag of the next runnable process identified at 14 is cleared at 20. If at 14 it was determined that the previous running process's flag had been set, then the method skips on to block 20 directly. At 22 the context switch to the new process is made and, at 24, the new process becomes the currently running process.

Back at block 14, should a next runnable process not have been found, then the method checks to see if the previous running process's flag has been set. If the flag has not been set, then the cached data is copied back at 28 and a flag in the previous running process's process descriptor is set at 30. A sleep mode is then entered at 32 until, for example, an interrupt occurs and the method returns to block 14 as part of an idle loop.

If at block 26 it is determined that the previous running process's flag has been set, then the method skips to block 30 as illustrated.

Fig. 2 illustrates computer system 34 embodying the present invention and which comprises the operating system kernel 36 and user mode processes 38. The user mode processes comprises three processes 40, 42 and 44 and the system illustrates the process 40 being blocked such that a scheduler 44 reschedules to process 42. Each of the runnable processes is

associated with a respective runnable process list 46, 48 having respective flags 46a, 48a. The configuration illustrated in Fig. 2 is such as to illustrate a context switch between processes and the manner in which the runnable process descriptor lists point to identify the runnable processes.

5 Lastly, Fig. 3 illustrates a CPU 50, data cache 52 and memory 54 of a computer system 56, and illustrates in particular the manner in which cached pages 58 – 62 are written back to respective memory locations 64 – 68 of the memory in accordance with the present invention. As will be appreciated, the data cache pages 58 - 62 are written back to the memory 54 if the process's
10 flag is not set.

Analysis on a Philips TriMedia 1300 (with an external memory manager unit) prototype board has shown that a typical 1Mb user mode application takes about 6000-7000 CPU cycles to write the process's data cache back to memory, which is about 30-50% of the total context switch time. With the data
15 cache being written to memory during CPU idle time, the context switch time loses this 30-50% overhead and so almost doubles the context switch speed. It will of course be appreciated that these figures vary, depending on the size of the application and the number of cached regions of memory.

The present invention finds particular use in UNIX-like systems, such as
20 Linux, FreeBSD, Solaris, etc and that use a processor with a level-1 or level-2 cache. All such UNIX-like systems have an idle loop, or idle process, that employ effectively wasted CPU cycles, and most modern processors have some form of memory cache in, or between, the processor and main memory.

The present invention can therefore find use in most current computer
25 operating systems.

As noted before, the invention can comprise a software, or hardware, solution to the problem of the prior art. If the idea is implemented in hardware, one of a variety of methods could be used. For example, the CPU sleep mode could be employed for the data cache write-back automatically during a sleep
30 period. Also, "flush data cache during sleep" flag in a status register could be employed to inform the CPU sleep operation to conduct the data cache write-back. Further, an extra sleep op-code could be added that performs the data

cache write-back, in addition to the normal sleep op-code. The addition of a data cache write-back operand to a CPU sleep opcode is also an option.

Thus, as will be appreciated, the present invention advantageously employs processor idle time to accelerate, and simplify, context switching in
5 computer operating systems. The invention overcomes the problem of relatively slow context switching speeds between runnable processes/tasks which arises due to the general requirement to write cached data back to memory. Since, in the prior-art, this writing-back of cached memory is carried out at the time of the actual context switch, the employment of otherwise spare
10 processor cycles in the idle loop of the operating system advantageously serves to enhance the operating speed since, at the time of initiating the context switch, the cached data has already been written back to memory.

CLAIMS

1. A method of context switching between processes in a computer operating system including writing cached data back to a memory means, comprising the step of the writing cached data back to the memory means
5 during processor idle cycles at completion of a process and prior to initiation of the context switch.

2. A method as claimed in Claim 1, including the step of setting a flag to indicate cached data has been written back to memory subsequent to
10 the completion of the process.

3. A method as claimed in Claim 2, including the step of, at the time of requiring a context switch, checking for the said flag to identify if the previous process's cached data has been written back to memory.
15

4. A method as claimed in Claim 3, including the step of, upon identifying that the said flag has been set, conducting the context switch without further cache write-back to memory.

20 5. A method as claimed in Claim 3, and including the step of conducting a cached data write-back to memory at the time of requiring a context switch if the said flag has not been identified as set.

25 6. A computer implemented system including a processor and cache memory arranged to receive operating system instructions for context switching between processes, and including control means for writing back cache data to memory means during processor idle cycle at completion of a process, and prior to initiation of the context switch.

30 7. A system as claimed in Claim 6, and arranged to employ a processor sleep mode to write cache data back to memory automatically during a sleep period.

8. A system as claimed in Claim 6, and including a status register having a flag for initiating data cache write-back during a CPU sleep operation.

5 9. A system as claimed in Claim 6, wherein the operating system is arranged with an additional sleep code instruction for the data cache write-back

10 10. A system as claimed in Claim 9, wherein the sleep code of the processor includes an additional data cache write-back instruction.

11. A computer program product having computer program instructions and arranged for controlling context switching between processors in a computer operating system so as to write cached data back to memory
15 means during processor idle cycles at completion of a process and prior to initiation of the context switch.

12. A computer program product as claimed in Claim 11, and including instructions for controlling to execution of a method as claimed in any
20 one or more of Claims 2 to 5.

1/2




FIG. 1




FIG. 2

2/2




FIG. 3

INTERNATIONAL SEARCH REPORT

PCT/IB 03/00626

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F9/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the International search (name of data base and, where practical, search terms used)

EPO-Internal, INSPEC, COMPENDEX, IBM-TDB, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>ALGUDADY M S ET AL: "A write update cache coherence protocol for MIN-based multiprocessors with accessibility-based split caches"</p> <p>PROCEEDINGS OF THE SUPERCOMPUTING CONFERENCE. NEW YORK, NOV. 12 - 16, 1990, WASHINGTON, IEEE COMP. SOC. PRESS, US, vol. CONF. 3, 12 November 1990 (1990-11-12), pages 544-553, XP010019975</p> <p>ISBN: 0-8186-2056-0</p> <p>page 545, right-hand column, line 24 - line 35</p> <p>-----</p> <p>-/--</p>	1-12

☒ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the International filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the International filing date but later than the priority date claimed

- *T* later document published after the International filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *G* document member of the same patent family

Date of the actual completion of the International search

20 December 2004

Date of mailing of the International search report

28/12/2004

Name and mailing address of the ISA
European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Müller, T

INTERNATIONAL SEARCH REPORT

PCT/IB 03/00626

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>ROTHMAN J B ET AL: "Sector cache design and performance" MODELING, ANALYSIS AND SIMULATION OF COMPUTER AND TELECOMMUNICATION SYSTEMS, 2000. PROCEEDINGS. 8TH INTERNATIONAL SYMPOSIUM ON SAN FRANCISCO, CA, USA 29 AUG.-1 SEPT. 2000, LOS ALAMITOS, CA, USA, IEEE COMPUT. SOC, US, 29 August 2000 (2000-08-29), pages 124-133, XP010515407 ISBN: 0-7695-0728-X page 125, right-hand column, line 46 - line 47</p>	1,6,11